

Effective networks for real-time distributed processing

Gonzalo Travieso* and Luciano da Fontoura Costa†

*Instituto de Física de São Carlos, Universidade de São Paulo,
Av. do Trabalhador São-Carlense 400, 13566-590, São Carlos, SP, Brazil*

The problem of real-time processing is one of the most challenging current issues in computer sciences. Because of the large amount of data to be treated in a limited period of time, parallel and distributed systems are required, whose performance depends on a series of factors including the interconnectivity of the processing elements, the application model and the communication protocol. Given their flexibility for representing and modeling natural and human-made systems (such as the Internet and WWW), complex networks have become a primary choice in many research areas. The current work presents how the concepts and methods of complex networks can be used to develop realistic models and simulations of distributed real-time system while taking into account two representative interconnection models: uniformly random and scale free (Barabási-Albert), including the presence of background traffic of messages. The interesting obtained results include the identification of the uniformly random interconnectivity scheme as being largely more efficient than the scale-free counterpart.

PACS numbers: 89.75.-k, 89.20.Ff

I. INTRODUCTION

We live in a world governed by action. From the ample motion of our planet to the intricacies of Brownian agitation, the universe is pervaded by an endless flow of changes to which our lives are no exception. While little can originate from stillness, movement imposes a continuing challenge to our senses. An immediate and important implication of movement is *causality*, one of the most essential elements in animal survival and also the key element in scientific investigation. In order to cope with such demands, animals evolved an intricate neuronal ‘hardware’ capable of analyzing moving images at a high resolution and rate appropriate to enable an immediate response, i.e. enough so as to favor their survival and reproduction. Such a type of reaction by dynamical systems is technically known as *real-time* (e.g., [1]). Despite the several advances in computing technology achieved along the last decades, we still lag well behind biological system as far as real-time processing and recognition is concerned. One of the possible ways to learn how to develop automated systems for effective, real-time processing is to look at the organization of biological systems for inspiration. Another possibility is to model and simulate such systems in order to try to identify particularly effective architectures and algorithms. One of the fundamental organizational principles of biological processing of information regards the inherent concurrency and parallelism characterizing those systems. Because neurons are relatively slow in processing and transmitting information (e.g., [2]), high speed can only be achieved by carefully interconnecting neurons so as to form groups or modules working in parallel. Indeed,

the brain is currently known to be organized according to interconnected modules [3] resembling a distributed computer system. In addition to the inherent features of the modules and involved neuronal cells, one particular feature of such modular processing systems concerns the specific way in which the several components are *interconnected*.

The interconnections between processing elements in a distributed system can be natural and effectively represented in terms of *complex networks* (e.g., [4, 5, 6]), where each processor is associated to a node while the interconnections between these nodes are expressed as edges. Through such a simple analogy, it is possible to bridge the gap between research in real-time distributed systems and the exciting concepts, tools and results from the area of complex networks. Although the origin of the latter area can be traced back to random graphs (e.g., [7]), and despite their immediate relationship with graph theory, the term *complex network* has been used to express the emphasis placed on graphs which exhibit complex structured connectivity.[19]

Thanks to technological advances in neuroanatomy and physiology, a more comprehensive vision of neuronal interconnections underlying the nervous systems of several animals is progressively emerging. At the microscopic — cellular — level, recent investigations have suggested that neurons are interconnected through small world and even scale free networks [8]. The macroscopic organization of cortical areas [3] also seems to be organized according to this principle [9].

This article presents the application of complex networks as the means for investigating the effect of alternative connectivity schemes, namely uniformly random (i.e. Erdős and Rényi — ER) and scale free (i.e. Barabási-Albert — BA) network models, on the overall performance of a real-time distributed processing system. While the ER model represents the natural reference system for connectivity, being almost universally considered

*Electronic address: gonzalo@ifsc.usp.br

†Electronic address: luciano@ifsc.usp.br

as the null hypothesis in complex network studies, the BA model is particularly representative of natural — including neuronal information processing systems [8] — and human-made systems such as the Internet and the WWW [4]. ER networks exhibit a characteristic node degree (i.e. the number of connections of each node of the network), in the sense that their overall connectivity can be well characterized in terms of the mean node degree. Contrariwise, BA networks exhibit a power law distribution of node degrees, which favors a heterogeneous connectivity, as well as the appearance of hubs (e.g., [4, 5]). In addition to their particular importance in modeling natural and human-made systems, BA networks provide an interesting model for the Internet and, consequently, grid computing systems — an important current trend in distributed computing [10] which provides a good deal of the motivation for the present work. In addition, the consideration of the BA model allows us to investigate the effect of the presence of hubs in parallel and grid systems — as implied by the Internet connectivity [11] — on the overall performance, as a counterpart to the otherwise almost regular connectivity ensured by the ER model making it similar to uniform parallel systems such as those involving mesh or hypercube interconnectivity. However, unlike those architectures, ER (and BA) networks are small-world.

The application model simulated here while considering these two interconnecting models involves a master node which distributes tasks, namely a stream of frames to be processed, among processing elements in other nodes acting as clients according to their availability. This assumes that both the source and destination of the frames are at the same site. The processing protocol considers a communication model involving routers connecting the clients to the master, as typically found in practice. Therefore, the overall modeling and simulation approaches adopted in this work include many realistic elements common to a real distributed processing system.

A previous work [12] studied the effect of interconnection topology in the performance of a grid computing application. The application considered in that work involved the processing of a number of not interacting tasks, with no real-time requirements. The lack of real-time constraints reduces the importance of traffic fluctuations enabling the use of average communication times for performance evaluation. Correspondingly, [12] does not include traffic effects. Under the real-time constraints of the application studied in the present work, varying delays induced by traffic play a major role, and application-independent traffic is thus included in the simulations. Another reason for inclusion of traffic is that the network topology strongly influences packet transit times under traffic, as reported in many works (e.g., [13, 14, 15, 16]).

The current article starts by presenting the adopted network, application and communication models, and follows by presenting and discussing the obtained results.

II. MODELS

The model used for the simulations comprises three components: a network model, a model for the communication between the collaborating computers, and an application model. These models are described below.

A. Network Models

When studying complex networks, the interest can be focused on their topologies, i.e. their structural properties, or on some dynamical processes taking place in the network. In this work, we use complex networks to describe the interconnection topology of a collection of computers participating in a collaborative real-time computation. As such, our main focus is on the dynamical processes of data communication and computations in the computers interconnected through the network. In this work, two widely used network models are considered: the Erdős-Rényi (ER) random network model with fixed number of edges [17] and the Barabási-Albert (BA) scale-free model [18].

ER networks are constructed by considering N isolated nodes (i.e. the network starts with no edges) and then adding edges one by one between uniformly chosen pairs of nodes (avoiding duplicate connections of nodes and self connections); the addition of edges is repeated a pre-specified number L of times. N and L are the parameters of the ER model and the average degree is given by:

$$\langle k \rangle = \frac{2L}{N}. \quad (1)$$

BA networks are constructed starting with m_0 nodes and adding new nodes one by one. When a new node is inserted, $m \leq m_0$ edges are added from this node to one of the previously existing nodes. The nodes to be linked are chosen following the *preferential attachment* rule [18]. The process is repeated until the desired number of nodes N is reached. In the simulations presented in this work, the initial network was fixed with $m_0 = 2m + 1$ fully connected nodes. The parameters of the model are thus N and m . Considering that for each new node m new edges are added and that the initial network already has m edges for each node, the total number of edges is $L = mN$, and therefore

$$\langle k \rangle = 2m. \quad (2)$$

B. Application Model

This work analyzes the influence of the network topology on real-time collaborative computation. The computation considered here is defined as follows: A special node in the network, called the *master*, is responsible for reading a stream of input data and writing a stream of

output data. The data arrives at the master in packets, here called *frames* in an analogy to real-time video processing, at regular intervals and the result of their processing must be output at the same interval.

For each input frame, an output frame is produced after the realization of a certain amount of computation (the computational load required for each frame). In this work, the load is considered equal for all frames. The computation is not done by the master. Instead, a collection of *clients* book their willingness to participate in the computation; when a new input frame arrives, the master chooses one available client and sends the frame to it for processing. After receiving and processing the input frame, the client sends the output frame back to the master; when the output frame arrives at the master, the client that processed the frame is again registered as ready to receive a new frame.

After arrival (or generation) at the master, each frame must be sent to a client, processed and sent back to the master. As communication delays in the network are unpredictable, the order of arrival of the resulting frames at the master is not guaranteed to correspond to the order in which they were originally delivered. To avoid output of the frames out-of-order and also enable waiting for the transmission and processing of the frames, a frame buffer must be maintained by the master, where arriving frames are stored in the correct order. The production of the output must then be delayed for some time, i.e. the output of frames must start some time after the arrival of the first frame. When a frame must be output, if it has not yet arrived it must be dropped with resulting quality loss. It is therefore important to allow sufficient time for the frames to arrive, but additional time given to frame processing then implies in increased latency in the production of the output. The time between the arrival of the first frame and the start of the output (which is also the time each frame will have available to be processed and returned to the master) is henceforth quantified in terms of the number of frame intervals.

C. Communication Model

After a network is generated according to a given model and set of parameters, its nodes are considered the *routers* of a computer network. The computers participating in the collaborative work are *hosts* connected to one of the routers. The master is connected to a router randomly selected with uniform probability. Not all of the network participates in the computation. The number of participating clients is a parameter of the simulation. Each client is associated with a router selected with uniform probability, but a limitation is imposed that each host (master or client) is associated with a different router. Only routers from the largest connected component of the network are selected.

As the network is assumed not to be exclusively dedicated to the frames computation, external traffic is sim-

ulated on the network by the generation of packets between random pairs of routers.

After insertion in the network, the packets are routed from node to node. The routers follow a “shortest path” routing strategy: each router sends a packet to a neighboring router that strictly decreases the number of steps remaining to reach the destination; if more than one neighbor satisfies this condition, one of them is chosen at random. While a router is routing and sending a packet to a neighbor, it cannot handle other packets. Packets arriving during this operation are queued in arrival order to be processed later; the routers are assumed to have unbounded queuing capacity.

The time for processing and communication at each step on the network is considered independent of the packet and router, although the delivery time for different packets might differ due to queuing. If the traffic in the network is low, the queues are empty or short, and the time taken for a packet to reach the destination is proportional to the topological distance between source and destination. As the traffic increases, congestion ensues [13, 15, 16], and the delivery time grows to many times that of the uncongested network.

D. Parameters

Here the model parameters and their values for the simulation results described below are presented.

Both network models are characterized by two parameters: the number of nodes N and the number of edges (for the ER model) or number of edges added for each new node (for the BA model). Henceforth the latter parameters are represented by the average node degree $\langle k \rangle$, that can be computed from the model parameters by using equations (1) and (2).

The computation dynamics is described by the computational load for the processing of each frame, the interval between frames and the number of frames to wait before starting the output. Considering that all clients are taken as identical (no difference in processing power), the computational load can be given as the computational time T of the processing task. The time interval between frames will be represented by τ and the number of frames to buffer by B . The output latency is therefore $B\tau$.

The time taken for a packet to traverse a step in the network from a node to one of its neighbors, h , is the same for each packet and router. As only the relation between the times are of importance, the time scale is chosen such that $h = 1$, the values of T and τ being expressed in these units. The random traffic generation in the network is assumed to be a Poisson process with inter-arrival times given by an exponential distribution with average $1/(N\lambda)$; the factor N is introduced to make the amount of traffic proportional to the size of the network; λ is the per-node packet generation frequency (in units compatible with $h = 1$).

The remaining parameter is simply the number of

TABLE I: Model parameters and their values. Time and frequency parameter “normalized” units (see text); output start interval in number of frame intervals.

Parameter	Meaning	Values
N	Number of nodes	1000
$\langle k \rangle$	Average node degree	2, 6, 10
T	Frame computation time	100
τ	Frame interval	5
B	Output start interval	10–50
λ	Packet generation frequency	0.001–0.02
C	Number of clients	100

clients $C \leq N - 1$. The parameters are listed in Table I, together with their range of values used in the simulations discussed below.

III. RESULTS AND DISCUSSION

A computation is successful if all output frames are returned from the clients and arrive at the master before they need for output. If a frame arrives too late for output, that frame is dropped, and the quality of the output is consequently reduced. Frames that arrive in time are here called *completed*. The number of completed frames is chosen as quality measure of the computation. In the simulations, a total of 1000 frames needs to be computed.

Figure 1 shows the number of completed frames as a function of network traffic and the output latency, for ER and BA networks of 1000 nodes, with $\langle k \rangle = 2, 6, 10$. The other simulation parameters are: 100 clients, frame interval of 5, frame processing time of 100. The results shown are averages of 100 simulations, each with a different network generated according to the corresponding model and different traffic patterns.

Consider first the case of the Erdős-Rényi network with $\langle k \rangle = 10$ (Fig. 1(e)). This plot shows a sharp transition on the number of completed frames for a latency of about 20 frame intervals. This transition is expected: with $T = 100$ and $\tau = 5$, at least $T/\tau = 20$ frame intervals must elapse before results start to arrive at the master. The fact that the transition is sharp, close to this lower limit, and independent of traffic in the studied region shows that an Erdős-Rényi topology with $\langle k \rangle = 10$ is efficient for this application, that is, it introduces small delays. For $\langle k \rangle = 6$ (Fig. 1(c)), the results are similar, but the transition is not so sharp and a larger latency is needed to reach the plateau of all frames completed. In the case of $\langle k \rangle = 2$ (Fig. 1(a)), another effect appears: a reduction on the number of completed frames occurs when the traffic is increased. The larger value of B needed and the drop in the number of completed

frames with increased traffic for reduced values of $\langle k \rangle$ are due to the reduction in the connectivity of the network: Few edges connecting the nodes result in increased average distances from the master to the clients; this affects the time taken to deliver the frames and complete their calculations, resulting in the need for an increase in the frame buffer and therefore larger latency. Also, the presence of fewer edges means that fewer alternative paths are available between the nodes, rising the sensitivity of the network to increased traffic.

For the Barabási-Albert networks, Figs. 1(b), (d), (f), the results show a much stronger influence of traffic. For $\langle k \rangle = 6$ and $\langle k \rangle = 10$ a continuous drop of the number of completed frames is noticed as the amount of traffic grows. For high traffic values, even large buffers are not able to guarantee the completion of a sufficient amount of frames. For $\langle k \rangle = 2$ the number of completed frames is small even for reduced amounts of traffic.

In order to better understand these results, Figure 2 plots the average packet transmission delay for the same situations as presented in Figure 1. The delay is computed as the time taken from the delivery of a packet at the source to the arrival at the destination. As shortest path routing is used and the time taken at each step (hop) is unitary, the average delay should equal the average distance between nodes under reduced traffic. This can be seen for the ER networks with $\langle k \rangle = 6$ and $\langle k \rangle = 10$ (Figs. 2(c),(e)), where the graphs are flat with a delay value about the value of the average distances. A different behavior is seen for ER networks with $\langle k \rangle = 2$ (Fig. 2(a)). At a packet generation frequency of about $\lambda = 0.01$, the delay starts to grow linearly with the amount of traffic. This is due to the onset of congestion in the network: some nodes start receiving packets far more frequently than they can handle, leading to increased queuing times of the packets in the nodes. After congestion, the average delays grow fast to many orders of magnitude of the average distance. Figure 2(f) shows that congestion occurs for the BA network with $\langle k \rangle = 10$ for a similar value of $\lambda = 0.01$, but note that the increase in delay is steeper after that point. For $\langle k \rangle = 6$, BA networks display congestion at lower traffics (about $\lambda = 0.005$) and even steeper increases of delay. The problem is accentuated for $\langle k \rangle = 2$ (Fig. 2(b)), where congestion occurs even for small amounts of traffic.

The reason for this greater sensitivity of the Barabási-Albert networks to traffic in comparison with the Erdős-Rényi counterparts can be easily understood. In fact, the preferential attachment rule of BA networks induces the creation of nodes with a high degree (hubs). Due to their high connectivity, these hubs appear in many of the shortest paths of the network. Although hubs are created, the number of hubs is always small, and most of the nodes have small connectivity and take part in just a few shortest paths. Therefore, a few nodes of the network become responsible for routing almost all of the traffic, resulting in large packet queues and congestion in these nodes. The lower the total connectivity of the network,

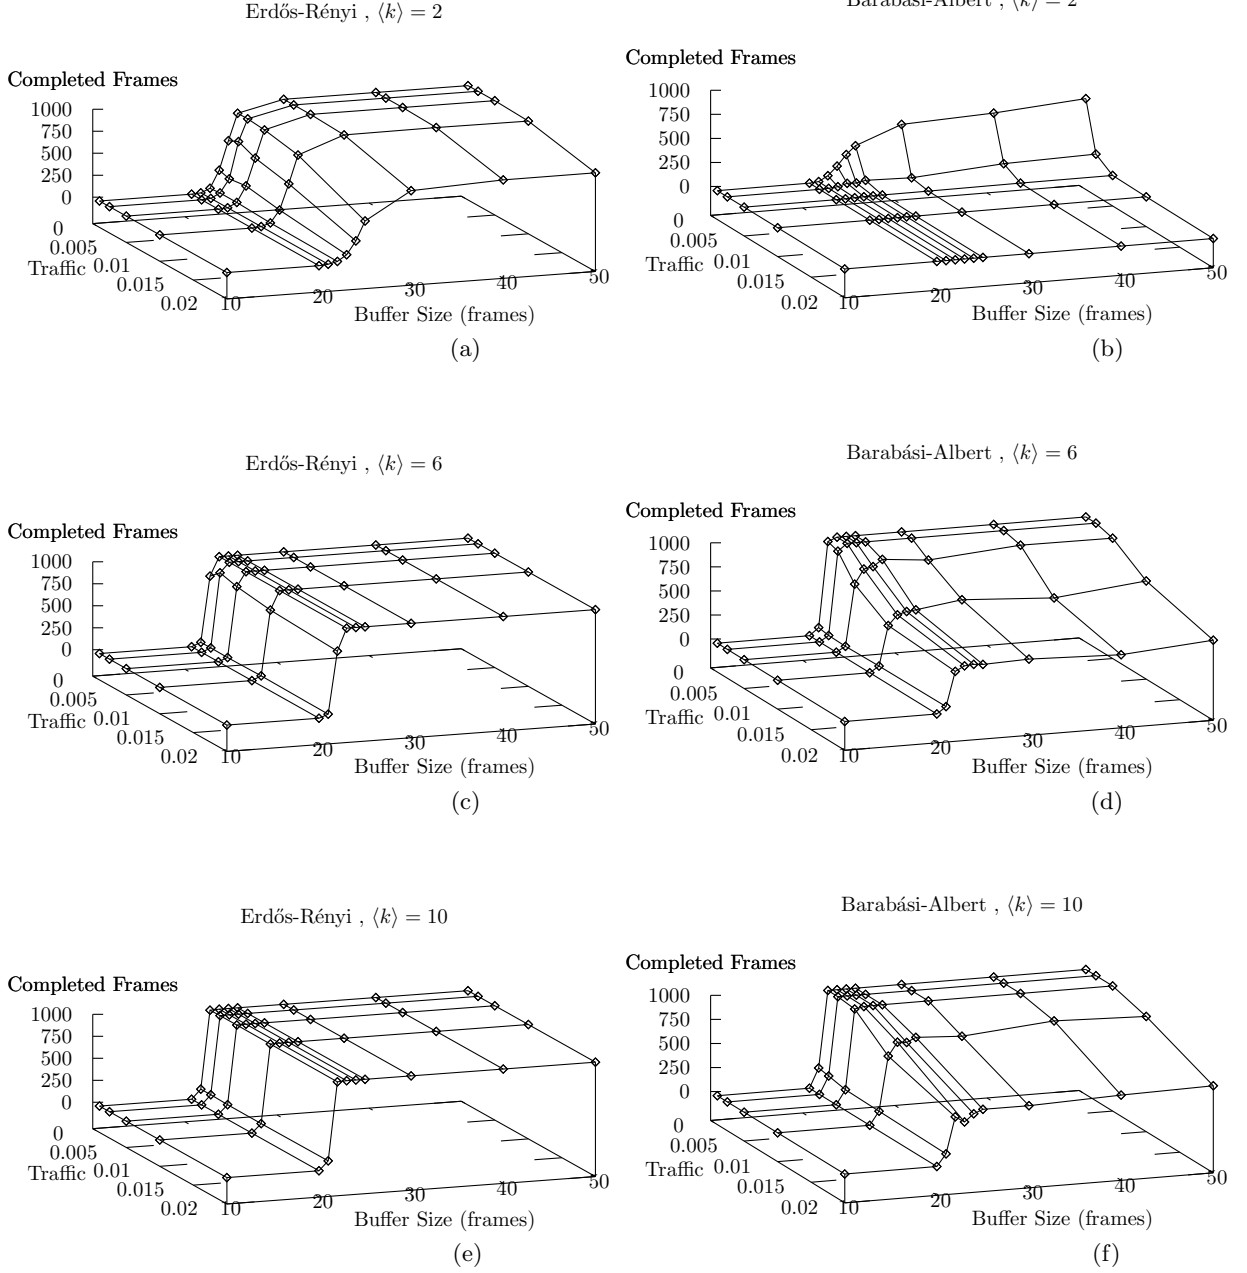


FIG. 1: Number of completed frames for a total of 1000 frames as a function of network traffic and number of buffered frames, for ER networks (left column) and BA networks (right column). Model parameters are $N = 1000$, $T = 100$, $\tau = 5$, and $C = 100$.

the more pronounced is this problem, as fewer links imply fewer alternative shortest paths. ER networks, on the other hand, distribute the connectivity homogeneously between all nodes, thus generating a better distribution of shortest paths among the nodes of the network.

To assess the influence of the computational load associated with each frame (parameter T), Figure 3 shows the number of completed frames as a function of traffic and frame processing time. For ER networks with $\langle k \rangle = 6$ and $\langle k \rangle = 10$, where no congestion occurs, two plateaux, one with all frames completed and the other with no frames completed, with a sharp transition be-

tween $T = 140$ and $T = 150$, are clearly seen. For the other cases, where traffic is important, a gradual decay of the number of completed frames is seen for increased traffic, as already seen in Figure 1, but there is also a gradual decrease of the number of completed frames as the frame computation time increases (before the transition to the no completion plateau). The higher the traffic, the steeper is the decrease of the number of completed frames with frame completion time.

The above results can be understood by the following reasoning. After the generation of a frame f , it must be delivered to a client, processed, and sent back to the

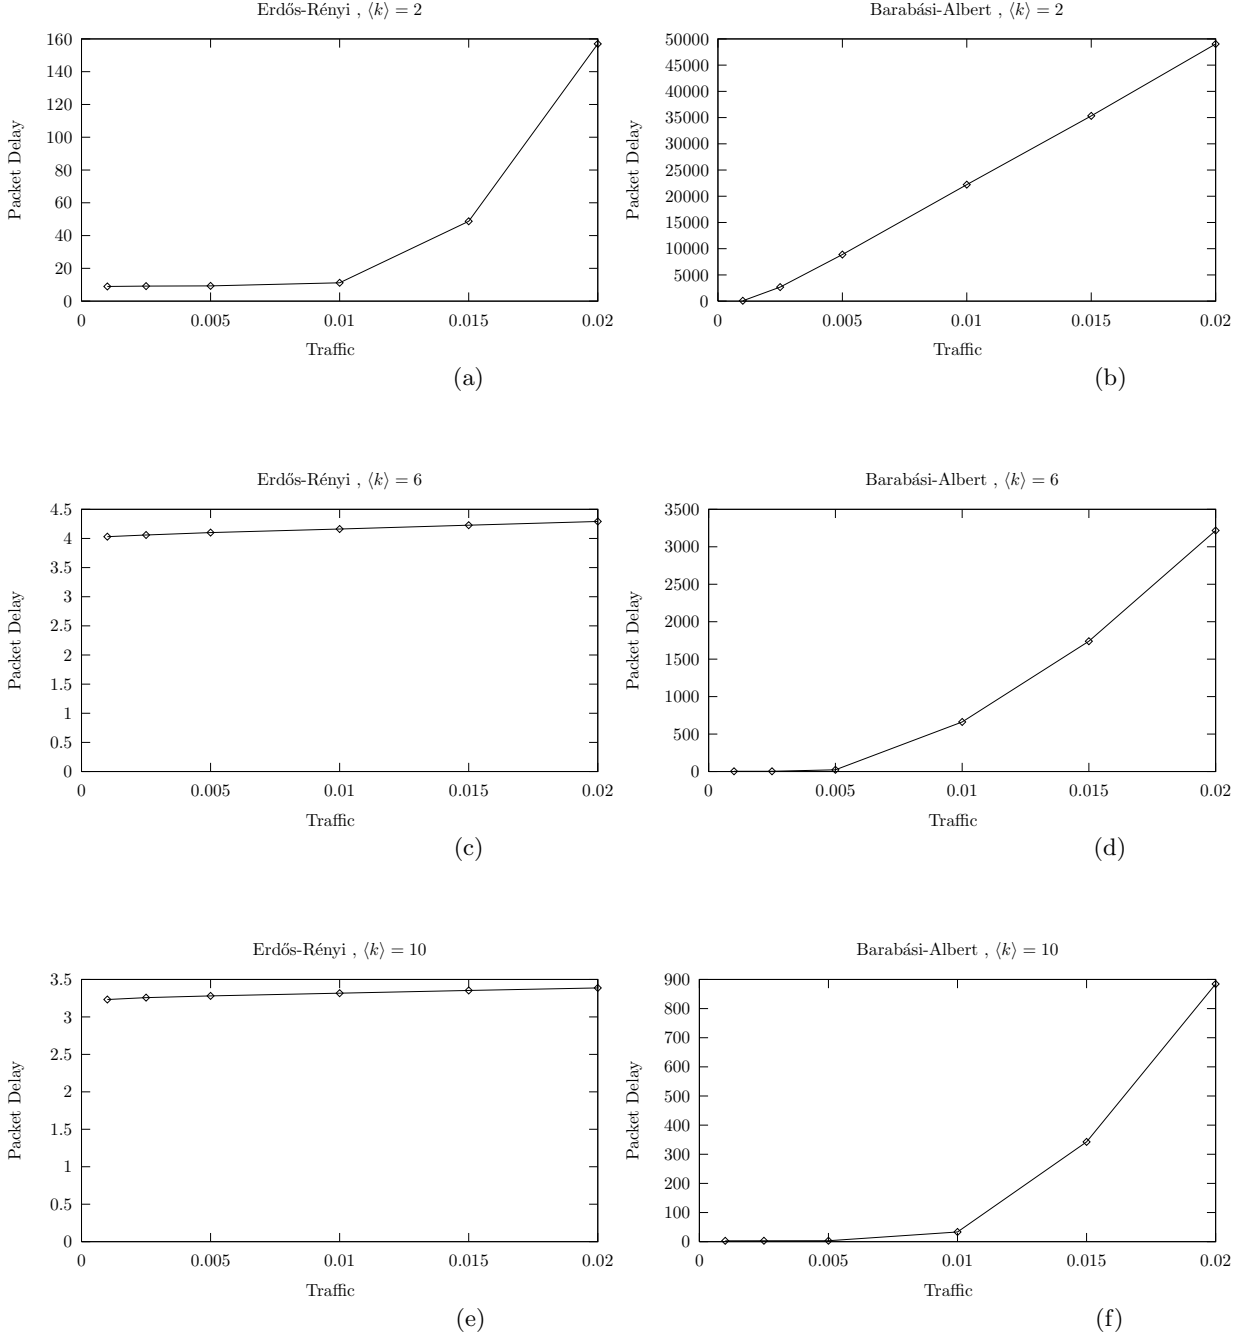


FIG. 2: Average delay for the delivery of packets in the network (time taken by the packets from source to destination) as a function of network traffic. All packets in the network are included in the average (not only packets that transport frames). Model parameters are $N = 1000$, $T = 100$, $B = 30$, $\tau = 5$, and $C = 100$.

master. Total processing time for f , $P(f)$ is given by

$$P(f) = w(f) + t_{mc}(f) + t_{cm}(f') + T \quad (3)$$

where $w(f)$ is the time f waits for a ready client, $t_{mc}(f)$ is the travel time of f from master to client, and $t_{cm}(f')$ is the travel time from client to master of the frame generated by the processing of f . Travel times $t_{mc}(f)$ and $t_{cm}(f')$ are generally different (although the topological distance is the same in both directions) due to possibly

different traffic conditions at the two transit periods. The condition for the completion in time of f is that $P(f)$ is less than the accepted latency $B\tau$, giving

$$w(f) + t_{mc}(f) + t_{cm}(f') + T \leq B\tau. \quad (4)$$

This condition must be satisfied by most frames. Under low traffic conditions, $t_{mc}(f) \approx t_{cm}(f')$ is close to the topological distance between master and client and small due to the small world property of the network models

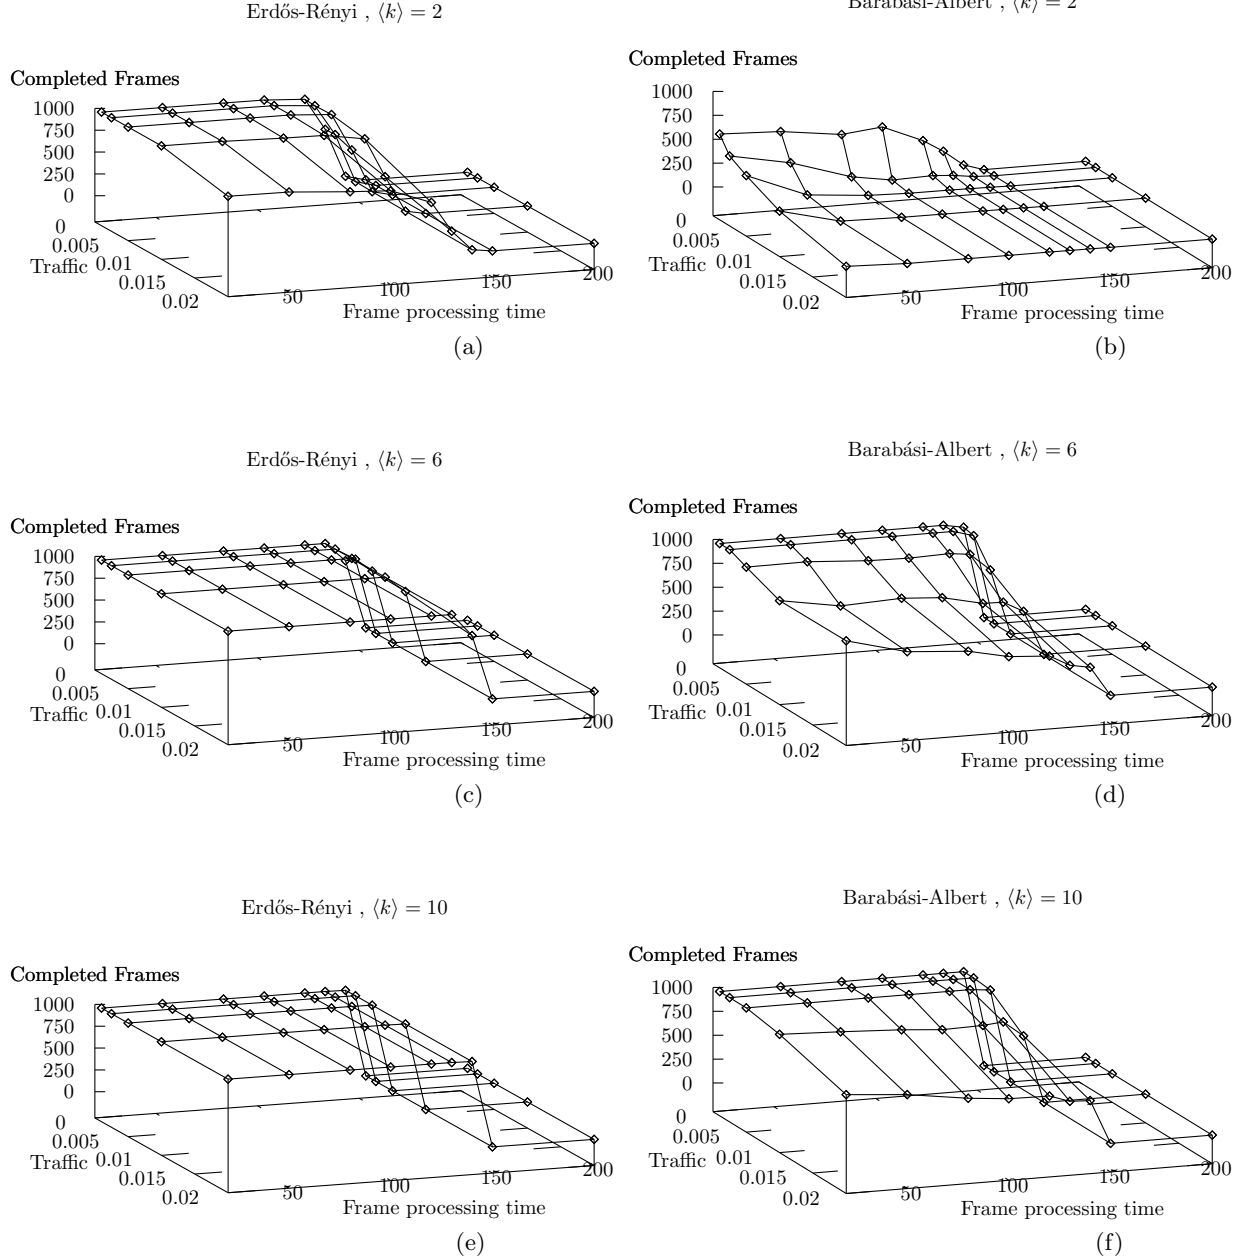


FIG. 3: Number of completed frames from a total of 1000 frames as a function of network traffic and frame computation time, for ER networks (left column) and BA networks (right column). Model parameters are $N = 1000$, $B = 30$, $\tau = 5$, and $C = 100$.

used, and the buffer used can be small, implying small latencies. Under heavy traffic, transit times can be very large (see Fig. 2), resulting in the need of high values of B and therefore large latencies; also, even with large B , the fluctuations in traffic are high, and many frames will be lost. This renders the distributed system useless for the application.

When a client returns the result of a computation to the master, it is automatically registered as able to receive a new packet. Therefore, it is reasonable to suppose that clients that communicate faster with the master will receive a larger number of frames to compute and as a

result the delays associated with the communication of frames can be smaller than the network averages. Also, as nearby clients tend to communicate faster with the master, the average distances (number of hops) traveled by frames can be lower than for the other packets. Figures. 4(a) and (b) show the distribution of packets with given per-hop delays (i.e., the ratio of the packet delay to the number of hops traversed by the packet) for the two considered network models and three different traffic conditions. It can be seen that under heavy traffic, specially for the Barabási-Albert model (which is more sensitive to traffic), the frames are subjected to smaller delays than

the other packets, but differences are substantial only for a small number of packets with prohibitively large delays, resulting in no advantage for the computation. The distribution of the number of hops traveled by all packet and the frames, for the two network models, are shown in Figures 4(c) and (d). They show that the clients effectively participating in the computation are uniformly distributed among the network nodes, with a perceptible change in distribution only for high traffic condition in the Barabási-Albert network model; in this latter case, effectively operating clients are positioned closer to the master node, implying that farther client nodes are receiving a smaller number of frames to be computed; again no advantage comes to the application being processed, as the difference occurs only in a traffic condition where the number of dropped frames is too high.

IV. CONCLUDING REMARKS

Combined with the availability of ever increasing amounts of data, the continuing advances scientific simulations have imposed serious demands for real-time processing. A natural means to cope with such a pressure is to develop and apply distributed systems, including the possibility of learning from biological systems and the application of Internet-based grid computing. Because of the high cost in implementing such solutions, it becomes essential to have access to realistic and effective modeling and simulation methodologies. The current work has described how concepts and methods from the modern area of complex networks research can be effectively applied in order to model and simulate with a good level of realism distributed systems for real-time processing, with emphasis focused on grid computing structures with connectivity underlined by the Internet. At the same time, because the BA model reflects some important connectivity features found in neuronal processing systems, the development and evaluation of such complex network mod-

els for real-time processing bear potential implications also for understanding biological processing.

Given its compatibility with some Internet topological features, and also because of its potential compatibility with neuronal processing systems, the Barabási-Albert complex network model has been selected in order to define the overall connectivity of the distributed real-time processing system. The Erdős-Rényi complex network model was also considered as a null hypothesis characterized by a high uniformity of node degree. Realistic models were assumed for the application and communication dynamics, including the effect of background message traffic, while the overall performance was quantified in terms of the total number of processed frames with respect to varying traffic intensity, buffer size and frame processing time. The obtained results included the identification of critical parameter configurations which are closely related to the model parameters and overall connectivity. Of special interest is the clear superiority of ER networks over BA networks. This is a result of the better handling of traffic by ER networks, as a consequence of the better distribution of connectivity between the nodes.

Possible future works include the consideration of other complex network models and applications, as well as inclusion of variability in the computing power of the clients, in the processing requirements of different frames and in the communication times between nodes. It is also of particular interest to study the effect of different routing algorithms and packet queuing strategies at the routers.

Acknowledgments

Luciano da F. Costa thanks FAPESP (05/00587-5) and CNPq (308231/03-1) for sponsorship. The work of G. Travieso was partially supported by FAPESP (03/08269-7).

-
- [1] J. W. S. Liu, *Real-Time Systems* (Prentice-Hall, 2000).
 - [2] E. R. Kandel, J. H. Schwartz, and T. M. Jessell, *Principles of Neural Science* (McGraw-Hill, 2000).
 - [3] S. Zeki, *Inner Vision: An Exploration of Art and the Brain* (Oxford University Press, 2000).
 - [4] R. Albert and A. L. Barabási, *Rev. Mod. Phys.* **74**, 47 (2002).
 - [5] M. E. J. Newman, *SIAM Review* **45**, 167 (2003), cond-mat/0303516.
 - [6] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang, *Physics Reports* (2006), accepted.
 - [7] B. Bollobás, *Random Graphs* (Cambridge University Press, 2001).
 - [8] O. Sporns, G. Tononi, and G. M. Edelman, *Cerebral Cortex* **10**, 127 (2000).
 - [9] O. Sporns, D. Chialvo, M. Kaiser, and C. C. Hilgetag, *Trends in Cognitive Sciences* **8**, 418 (2004).
 - [10] I. Foster, C. Kesselman, and S. Tuecke, *International Journal of High Performance Computing Applications* **15**, 200 (2001).
 - [11] M. Faloutsos, P. Faloutsos, and C. Faloutsos, in *SIGCOMM '99: Proceedings of the conference on applications, technologies, architectures, and protocols for computer communication* (1999), pp. 251–262.
 - [12] L. da F. Costa, G. Travieso, and C. A. Ruggiero, *European Physical Journal B* **44**, 119 (2005), cond-mat/0312603.
 - [13] P. Holme, *Advances in Complex Systems* **6**, 163 (2003).
 - [14] B. Tadić, S. Thurner, and G. J. Rodgers, *Physical Review E* **69**, 036102 (2004).
 - [15] L. Zhao, Y.-C. Lai, K. Park, and N. Ye, *Physical Review E* **71**, 026125 (2005).
 - [16] B. Tadić, G. J. Rodgers, and S. Thurner (2006), physics/0606166.

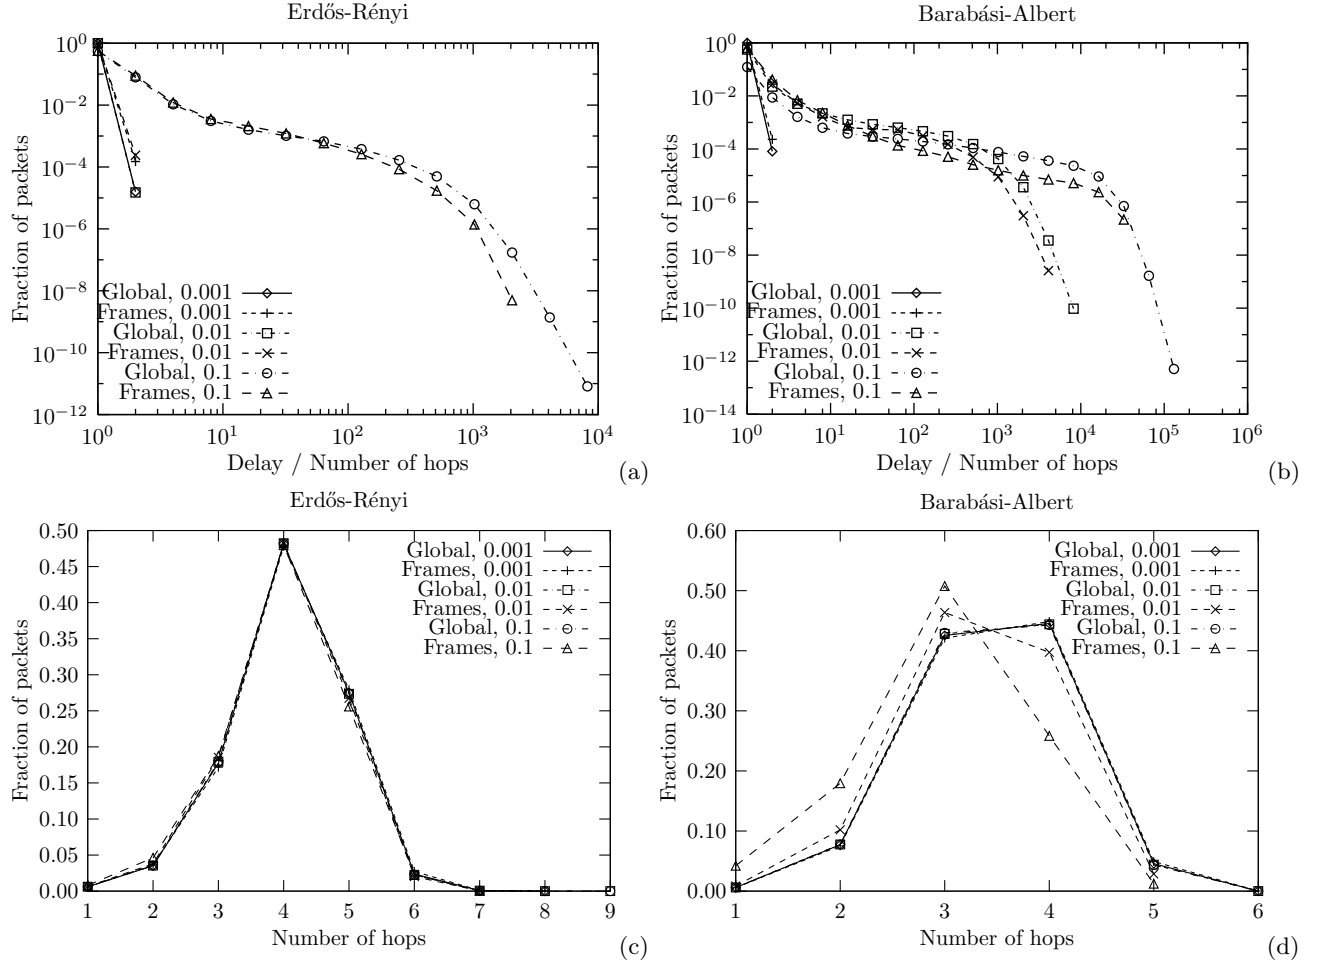


FIG. 4: Communication properties of frames (Frames) as compared to all packets (Global) in the network: per hop delays (a) and (b) and number of hops (c) and (d), for the Erdős-Rényi (a) and (c) and Barabási-Albert (b) and (d) network models. Simulation parameters are as in Table I and $\langle k \rangle = 6$, $B = 50$, and $\lambda = 0.001, 0.01, 0.1$.

- [17] P. Erdős and A. Rényi, *Publicationes Mathematicae* **6**, 290 (1959).
- [18] A.-L. Barabási and R. Albert, *Science* **286**, 509 (1997).
- [19] Although authors tend not to consider uniformly random graphs (constant probability of connection between any

pair of nodes) as complex networks, even such simpler structures do exhibit complex connections as a consequence of statistical fluctuations.